

DTIC FILE COPY

NUSC Reprint Report 6570
17 November 1986

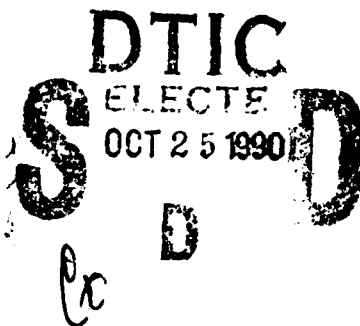
(1)

TARSIA: An Intelligent System for Underwater Tracking

J. G. Baylog
S. M. Zile
J. T. Rumbut

Combat Control Systems Department

AD-A227 792



Naval Underwater Systems Center
Newport, Rhode Island / New London, Connecticut

Approved for public release; distribution unlimited.

Reprint of a paper presented at the Expert Systems
in Government Conference, McLean, VA, October 1986.

90 10 22 067

TARSIA: AN INTELLIGENT SYSTEM FOR UNDERWATER TRACKING

John G. Baylog Suzanne M. Zile John T. Rumbut

Naval Underwater Systems Center
Newport, RI 02841-5047

ABSTRACT

This paper describes TARSIA, a knowledge-based system for underwater contact tracking. In particular, the tracking problems addressed by TARSIA encompass contact solution quality assessment and multiple solution integration and these are presented in a preliminary discussion. Subsequently, the function of TARSIA is described, domain specific knowledge representation requirements are formalized and TARSIA reasoning capabilities are summarized. Salient design features include temporal reasoning, symbolic/numeric interaction, and a mixed-initiative user interface.

Introduction

The TARSIA (TARgeting Solution Integration and Assessment) domain involves data and information processing for contact tracking in an underwater acoustic environment. Signals either emitted from or reflected off of a contact are received by an observer and processed to provide measurement data such as the azimuthal angle of arrival (bearing), indicated in figure 1. As measurement data is collected over time, it is further processed to generate a contact solution, that is, a set of parameters related to contact state (e.g., position and velocity). However, adverse tracking conditions are often encountered in the underwater acoustic environment. Measurement data is typically corrupted by additive noise indicative of ambient noise in the sound channel and by systematic errors that arise during data generation. Consequently, statistical estimation techniques are utilized to minimize error in contact state estimates and to provide an indication of solution accuracy, referred to as solution quality.

Measurement data may be available from several sources (e.g., different sensors deployed by the observer), and when processed independently, yield multiple solutions. The need exists to combine these local solutions into an optimal global solution upon which command level decisions can be based. Successful multiple solution integration, however, requires accurate solution quality assessment of the component local solutions, or the resulting global solution and solution quality will be misleading¹. In practice, encumbrances due to data anomalies or mismodelling

inherent in the estimation process are encountered and oftentimes are not properly taken into account, resulting in solution integration and assessment inaccuracies.

Conventional tracking technology has traditionally taken an algorithmic approach. However, the integration and assessment of solutions is knowledge-intensive, often requiring a user to perform these tasks based on his experiential knowledge. Complexity of solution integration and assessment is likely to multiply as advances in contact tracking technology occur, exposing personnel to increasing amounts of information to assimilate. In view of these facts, the development of TARSIA was undertaken to explore a knowledgebased approach to solution integration and assessment². What follows is an overview of the relevant aspects of contact tracking, and a discussion of the functional goals of TARSIA. The knowledge engineering of TARSIA is subsequently described, along with pragmatic issues confronted during the system design.

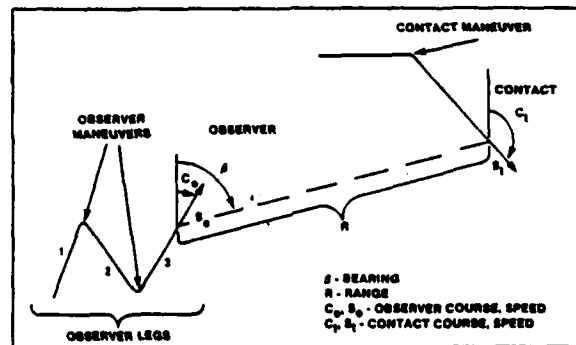


Figure 1. Contact - Observer Geometry

Contact Tracking Requirements

Various estimation techniques are used to process measurement data to provide contact solutions. The ability of an estimator to extract desired contact parameters from measurement data is dependent upon factors such as parameter observability, contact motion modelling, and characterization of the acoustic environment³. Param-

Availability Codes	
Dist	Avail and/or Special
A-1	

eter observability represents the most basic requirement for contact tracking and indicates the kind of information that can be extracted from a sequence of measurements of specified type. For example, when bearing measurements become available to a mobile observer travelling at constant velocity, a contact range estimate is unavailable until the observer maneuvers to effect a change in velocity⁴.

While observer motion is generally considered unrestricted, more restrictive models of contact motion are utilized to expedite the estimation process. Typically, a constant velocity model is used whereby the estimated contact trajectory is comprised of piecewise linear segments, referred to as legs, over which the constant velocity assumption is valid. Any maneuvers performed by the contact violate this assumption and provision must be made for the determination of boundaries in the observed data over which constant contact velocity can be presumed.

The acoustic environment plays a dominant role in the availability and interpretation of measurement data. Variations in the speed of sound propagation occur due to changes in water temperature, salinity, and pressure, causing the refraction of sound waves as they propagate through the medium. Reflections from the ocean surface, bottom, or existent thermoclines further complicate signal and data processing by causing multiple sound paths to exist concurrently. Variations in boundary conditions such as sea state, and bottom depth, slope, and composition can distort signals or preclude accurate interpretation of data derived from them. When multiple spatially-separated sensors are deployed by the observer, sound propagation from the contact to each sensor must also be considered.

Estimation Techniques

Parameter estimation in an underwater contact tracking environment often involves the exploitation of nonlinear relationships between measurement data and the parameters of interest. Computer-based nonlinear estimation techniques successfully applied to underwater contact tracking include computer-operator interactive, batch processing, and sequential processing techniques. Each of these involves varying levels of operator interaction as indicated in table 1.

An underlying component of each of these estimation methods is that of measurement residuals. A residual sequence is defined as the difference between the available measurements and their predicted values based on current solution estimates. Estimation algorithms are often designed to minimize the resulting set of measurement residuals. In effect, this results in the extraction of contact related information from the residual sequence. Data and/or modelling anomalies can cause patterns to remain in a residual sequence subsequent to minimization, thus providing a source of evidence for determining the existence and nature of mismodelling.

When multiple data sources are available, optimal estimation criteria would dictate processing all the information to achieve an optimal solution with the smallest possible uncertainty region. However, modelling assumptions inherent in the estimation process may be inadequate or inaccurate, and generating an optimal global solution free of modelling error may not be possible. An alternative is to develop multiple local solutions where data consistency is evident⁵; this provides a mechanism for determining the relative quality of various data sets. Subsequently, knowledge of local solutions can be combined with ancillary knowledge of the acoustic or tactical environment to form the best global estimate of contact state.

The Goals of TARSIA

In determining the specific function of TARSIA, focus has been placed on solution quality assessment and multiple solution integration. In particular, the functional goals of TARSIA include:

1. Select contact parameters to be estimated and data subsets to be processed;
2. Monitor the data processing to ascertain whether the data supports the modelling assumptions;
3. Interpret the quality of local and global solutions and assess the status of contact tracking;
4. Control the solution integration process.

Solution integration in TARSIA entails solution selection, solution quality modification, and the utilization of modelling assumptions, hypothesis testing procedures, and experiential knowledge.

Table 1. Estimation Techniques

ESTIMATION METHODS	DATA SET SELECTION	PARAMETER ESTIMATION	COMMENTS
COMPUTER-OPERATOR INTERACTIVE	MANUAL	MANUAL	CAN BE OPERATOR INTENSIVE OPERATES ON SPECIFIED DATA SET
BATCH	AUTOMATIC OR MANUAL	AUTOMATIC	CAN BE COMPUTATIONALLY EXPENSIVE OPERATES ON SPECIFIED DATA SET
SEQUENTIAL	AUTOMATIC	AUTOMATIC	COMPUTATIONALLY EFFICIENT OPERATES ONLY ON MOST RECENT DATA

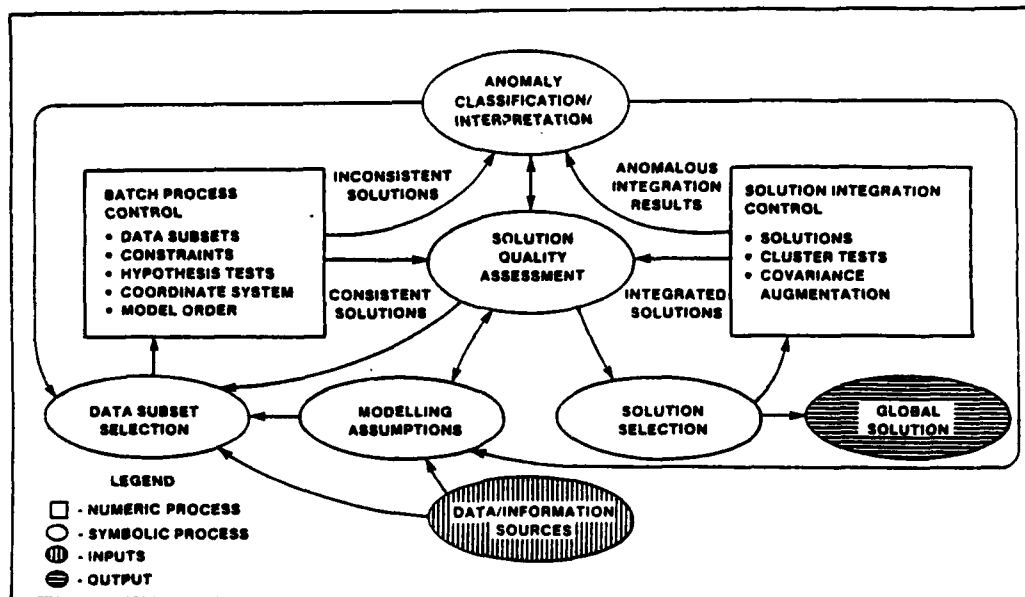


Figure 2. Conceptual Model of TARSIA

Estimation techniques are intrinsic to the TARSIA domain, and of the various methods considered, batch processing was selected as the primary estimation tool. This approach was selected to allow for reasoning over sets of data while retaining the algorithmic flexibility to automatically formulate solutions. Intelligent control of the batch processor entails utilization of many of the heuristics inherent in the computer-operator interactive techniques. The batch processor, however, imposes less of a burden on the system user and is more amenable to direct accumulation of evidence regarding the status of contact tracking. Furthermore, constrained estimation techniques⁶ can be employed to utilize a priori information and allow for division of the parameter search space. Sequential estimators were not considered since their performance can degrade due to linearization error when operating under marginal conditions⁷.

With the TARSIA functional goals firmly established, a conceptual system was outlined and is shown in figure 2. In this conceptual model, interaction between numeric and symbolic components is triggered by the reception of knowledge (e.g., the availability of measurement data). Invocation of numeric processes is controlled symbolically, where for each data subset processed, the batch processor yields individual solutions as well as an associated solution quality. Successive iteration between TARSIA components yields a global solution with its interpreted solution quality.

Types of Knowledge

The knowledge space of TARSIA consists of multiple knowledge types with both perceptual and abstract components. Perceptual knowledge is a

product of sensory input (e.g., acoustic and non-acoustic data), while abstract knowledge encompasses the understanding of physical processes and tracking phenomena (e.g., parameter observability). As the granularity of perceptual knowledge is generally coarse, it undergoes refinement prior to use with abstract knowledge. The knowledge space is further qualified from a representation and utilization perspective, where knowledge types are characterized as: factual, temporal, procedural, evidential and heuristic. Heuristic knowledge combines experiential knowledge together with the aforementioned knowledge types and is used to provide various levels of decision making. In the TARSIA domain, heuristic knowledge is diverse, ranging from the criteria for data subset pruning to methods for determining whether an anomaly exists. Examples of TARSIA knowledge types are provided in table 2.

Factual knowledge in the TARSIA knowledge base includes (but is not limited to) general and specific facts about the acoustic environment, the tactical environment, sensors, acoustic and non-acoustic data, data anomalies, parameter observability, modelling assumptions, estimation techniques, and solution information. Temporal knowledge, a form of factual knowledge, is an underlying component of the TARSIA domain. Tracking events occur, data is lost and regained, and solutions must be developed based on past information. As such, TARSIA must be able to quantify time, identify and reason about temporal trends in events and/or data, and understand symbolic concepts such as before, during, and after.

Procedural knowledge involves the step-wise application of numeric or symbolic operations to obtain a specified goal. An example of procedural

Table 2. TARSIA Knowledge Types and Examples

FACTUAL	<ul style="list-style-type: none"> • AN ACOUSTIC CONVERGENCE ZONE EXISTS WITH ANNULUS BETWEEN R_{MIN} AND R_{MAX} NAUTICAL MILES. • CONTACT 1 HAS A MAXIMUM SPEED OF S_{MAX} KNOTS.
TEMPORAL	<ul style="list-style-type: none"> • BEARING WAS AVAILABLE FROM 1023 THROUGH 1053. • BEARING WAS AVAILABLE DURING THE LAST SEGMENT IN WHICH THE OBSERVER TRANSITED THROUGH TWO LEGS.
PROCEDURAL	<ul style="list-style-type: none"> • DETERMINE ALL POSSIBLE MEASUREMENT PARTITIONS, PRUNE PARTITIONS TO A MANAGEABLE NUMBER, AND SEND PARTITIONS TO THE BATCH PROCESSOR.
EVIDENTIAL	<ul style="list-style-type: none"> • BIAS WAVEFORMS EXHIBITED IN THE TIME DEALY RESIDUALS FOR SEVERAL CONTACTS INDICATE THAT SENSOR ORIENTATION IS SUSPECT.
HEURISTIC	<ul style="list-style-type: none"> • IF BOTTOM BOUNCE RELATED MEASUREMENTS EXPERIENCE ANOMALIES, THEN LOOK FOR MODELLING ERROR IN PROPAGATION MODE OR BOTTOM ATTRIBUTES.

knowledge is the selection of measurement data subsets over which to form local solutions via the batch processor. This is accomplished through data partitioning and/or segmentation and identification of the most appropriate subsets for solution generation. In partitioning, sets of measurement data types are formed by utilizing parameter observability and a selected set of modelling assumptions as the partitioning metrics, whereas segmentation involves the temporal division of measurement data. Figure 3 illustrates subsets formulated as a result of partitioning and segmentation, but more typically, subsets are composed utilizing both processes. Following subset formulation, pruning operations yield the most appropriate subsets.

PARTITIONING	
SUBSET 1 BEARING	(0630-0700)
SUBSET 2 CONICAL ANGLE	(0630-0700)
SEGMENTATION	
SUBSET 1 BEARING	(0630-0645)
CONICAL ANGLE	(0630-0645)
SUBSET 2 BEARING	(0645-0700)
CONICAL BEARING	(0645-0700)

Figure 3. Data Partitioning and Segmentation

Evidential knowledge is derived from factual features extracted from measurement data, observer trajectory, sensor orientation, and other information sources, and is primarily used for anomaly classification and interpretation. Possible causes of anomalies can be grouped into a frame of discernment⁸ and include measurement anomalies (e.g., bias, wild points, and tracker errors), contact motion mismodelling, sensor position and orientation mismodelling, environmental anomalies (e.g., propagation path mismatches, boundary mismodelling, unmodelled refraction, and interference), and contact/data association errors. Given

a data set and its resulting solution, confidence in the solution quality is established if the measurement residuals are deemed free of anomalies. Conversely, bias detected in the residuals indicates the existence of an anomaly, where the form of the bias lends support to certain hypotheses within the frame of discernment.

Knowledge Representation

A hierarchical frame-based representation was selected as the most suitable for TARSIA, since much of the knowledge was amenable to generalization and classification. This representation enables use of inherited default values in cases where uncertainty exists but limiting knowledge is helpful. Though the TARSIA knowledge base contains many frames, the five generic frames used most often include the measurement, observer, contact, sensor, and estimator frames. Other frames may be added as TARSIA evolves.

As mentioned earlier, time must be accounted for in the TARSIA system. Various temporal representation methods were reviewed for their suitability to the domain. The notion of reference events⁹ has been integrated with an interval-based (versus point-based) representation, similar to the work of Allen and Vilain^{10,11}, for representation of temporal knowledge. This representation has been incorporated into the overall TARSIA frame/slot representation.

Used as a method of organizing facts and providing constraint propagation, reference events are occurrences that have a well-known time and are often used to assess the times of events related to them. Reference events are chosen to "correspond to key events that naturally divide the facts in the domain". Examples of TARSIA reference events are observer maneuvers and legs, contact maneuvers and legs, and measurement availability fluctuations. Slots, such as data time in figure 4, are reference event slots whose values are time intervals over which the reference events have occurred. Often, slots in the same frame access the values of reference event slots via time tags since the temporal information is the same.

BEARING FRAME

CONTACT NUMBER: 1
DATA TIME: (2 (0645 CURRENT)) (1 (0645 0645))
DATA ENTRY MODE: (2 CONTINUOUS) (1 INTERMITTENT)
DATA QUALITY: (FAIR (0657 CURRENT))
(GOOD (0645 0657)) (FAIR 0635 0645))

Figure 4. Example of Temporal Interval Representation

Various slot representations are required for: 1) time independent information that requires no temporal context; 2) reference event slots; 3) slots that access the information found in reference event slots; and 4) slots whose temporal context varies independent of the reference event. These have been implemented as: 1) a value; 2) a list composed of a time tag and a time interval; 3) a list composed of a time tag and a value; or 4) a list composed of a value and its associated interval of time.

TARSIA represents the present as "current", where any interval with a "current" time is on-going, thus facilitating the notion of persistence¹². Time in TARSIA is represented in seconds, although it is intended that the system handle other units of time as well. Having restricted TARSIA to events with well established times, it does not deal with fuzziness of event occurrence.

Evidential knowledge comprises that portion of the TARSIA knowledge base used for the classification of detected anomalies. Shafer's representation¹³ has been selected for implementation because ignorance of the nature of an anomaly can be explicitly represented. This representation expresses the belief in a hypothesis X by an interval $[s(X), p(X)]$, where $s(X)$ represents support for hypothesis X and $p(X)$ denotes its plausibility. Each hypothesis X is represented as a frame, and in addition to support and plausibility, contains slots for basic probability assignment, temporal context, evidence, and evidence sources.

Reasoning and Control

Given its diverse types of knowledge, several types of interrelated reasoning capabilities are required by TARSIA to perform decision making. Temporal reasoning is integral to each of these, providing for the generation and utilization of historical knowledge concerning past events as a basis for problem solving. Other reasoning types in TARSIA include system control reasoning, data/event specific reasoning, and evidential reasoning.

System control reasoning involves meta-knowledge pertaining to overall contact tracking objectives, decision-making criteria, and control of algorithmic functions. It also includes knowledge of alternate decisions and the consequences of making these decisions. Finally, control reasoning establishes when and where operator-supplied information and control occurs.

Data/event specific reasoning in TARSIA is employed by system control whenever intermediate decisions must be made based on the occurrence of tracking events. Data specific reasoning can occur in response to measurement "events", such as the availability of measurement data, where utilization decisions must be made. As tracking events take place (e.g., the execution of observer or contact maneuvers), event specific reasoning determines the effect of these events on parameter observability, solution quality and solution integration.

Evidential reasoning, a form of data/event specific reasoning, is used to determine the nature of data anomalies and ascertain the validity of modelling assumptions. Inconsistencies in measurement residuals or between solutions, are detected via numeric procedures such as hypothesis testing and clustering techniques or through operator specification. Attributes of the anomaly (e.g., temporal duration) are then collected which are used to govern the distribution of belief amongst the various hypotheses in the frame of discernment. Anomaly classification, if possible, is subsequently performed based on the strength of belief associated with the related hypotheses.

To illustrate the interaction between reasoning components, the following example is presented. Consider the case where mismodelling has occurred due to the existence of an undetected contact maneuver within a specified data subset. Following subset selection, a solution is generated by the batch processor. If conditions permit (e.g., adequate data quality and contact-observer geometry), TARSIA detects the existence of an anomaly in the measurement residuals. Given this evidence, data sifting procedures are employed and the results interpreted by TARSIA to infer the temporal boundaries and duration of the anomaly. Existence of a short duration anomaly is deduced and the data is segmented at the anomaly boundaries. From this, support is assigned to the hypothesis of a contact maneuver and the batch processor is invoked to operate over the segmented intervals, yielding solutions for each. Further evidence is provided to TARSIA when, over the respective intervals, consistency of the residuals is detected and distinct velocity estimates of sufficient quality are produced. When this evidence is combined, the contact maneuver hypothesis is asserted as being true.

System Architecture

TARSIA is being developed using both an LML Lambda 2X2+ and a VAX 11/780, with the knowledge-based component of TARSIA on the Lambda and much of the numeric processing capabilities on the VAX. Additionally, the user will be provided with a capability to visually inspect information produced as a result of numeric processing (e.g., measurement residuals). IntelliCorp's Knowledge Engineering Environment (KEE) is used to represent the knowledge of TARSIA. Forward chaining rules, implemented in Lisp, provide the reasoning mechanism of TARSIA and are accessed either programmatically or through demons.

Table 3. Packet Types and Characteristics

PACKET NAME	COMMUNICATION FLOW	SENDING CRITERIA	PACKET CONTENTS
SCENARIO STARTUP (SSP)	VAX-LAMBDA	INITIALIZATION OF SCENARIO	INITIAL OBSERVER PARAMETERS ENVIRONMENTAL CONDITIONS NUMBER OF KNOWN CONTACTS SENSORS DEPLOYED
OBSERVER TRACK (OTP)	VAX-LAMBDA	OBSERVER EVENTS- LEG, MANEUVER, SENSOR STABILIZATION	OBSERVER PARAMETERS OBSERVER MOTION GOALS SENSOR ATTRIBUTES
MEASUREMENT AVAILABILITY (MAP)	VAX-LAMBDA	MEASUREMENT EVENTS- NEW, LOST, REGAINED, ATTRIBUTE CHANGE	CONTACT NUMBER SENSOR SOURCE MEASUREMENT ATTRIBUTES PROPAGATION PATH
BATCH INSTRUCTION (BIP)	LAMBDA-VAX	WHEN SOLUTIONS ARE REQUIRED POSSIBLY OPERATOR ACTIVATED	CONTACT NUMBER MODEL ORDER DATA SUBSETS PARAMETER CONSTRAINTS INITIALIZATION METHOD HYPOTHESIS TEST REQUEST
SOLUTION INFORMATION (SIP)	VAX-LAMBDA	IN RESPONSE TO BIP	CONTACT SOLUTION SOLUTION QUALITY HYPOTHESIS TEST RESULTS
SOLUTION INTEGRATION INSTRUCTION (SIIP)	LAMBDA-VAX	WHEN SOLUTION INTE- GRATION IS REQUIRED POSSIBLY OPERATOR ACTIVATED	CONTACT NUMBER SOLUTIONS CLUSTER TEST REQUESTS COVARIANCE AUGMENTATION
INTEGRATED SOLUTION INFORMATION (ISIP)	VAX-LAMBDA	IN RESPONSE TO SIIP	INTEGRATED SOLUTIONS SOLUTION QUALITY CLUSTER TEST RESULTS

Symbolic/Numeric Communications

Communication is accomplished via sequential message passing between the VAX and Lambda in the form of information packets. Several generic packet types have been developed to transfer different kinds of information; these packets and their associated characteristics are described in table 3. Packets generated on the VAX are classified as information packets, whereas packets produced by the Lambda are viewed as control packets. From the VAX, the occurrence of tracking events cause the generation and sending of packets to the Lambda. As these packets are processed, new or updated information becomes available to TARSIA, instances of frames are created, slot values are updated, demons are activated and rule clusters are fired. Control packet generation is initiated as TARSIA requires certain types of numeric processing to be performed (such as solution integration).

Mixed-Initiative User Interface

In certain situations the operator may forego contributing to the solution selection process. However, during normal operation his insights can be integrated into the TARSIA knowledge base. Operator specified information is made available to TARSIA through a mixed-initiative user interface; this interface will ultimately provide an explanation capability and a menu format query capability. By making the operator an active

participant¹⁴, TARSIA can provide informative counsel throughout the solution integration and assessment process, in addition to an integrated global solution.

Possible levels of operator participation include automatic, limited interaction, and extended interaction. In the automatic mode, TARSIA operates as a background process and relies only on data available from its sensors and the operator's decision making preferences (e.g., when automatic solution generation should occur), established prior to scenario commencement. The limited interaction mode will allow the operator to specify certain kinds of information that is not currently available to the knowledge-based system. In this mode, interim explanations would be available at prespecified points in TARSIA processing. The extended interaction mode will allow the operator to enumerate specific functions that he wishes to effect and will include a limited vocabulary query capability.

Summary and Conclusions

TARSIA is a knowledge-based system designed to perform solution integration and quality assessment in an environment where a propensity for modelling error exists. TARSIA was implemented utilizing a hierarchical frame-based representation where reasoning is accomplished via forward chaining rules.

The feasibility of applying Artificial Intelligence technology to the contact tracking problem has been demonstrated. In doing so, several germane issues were resolved. In particular, the foundation for temporal reasoning in the TARSIA domain has been established and is applicable to other underwater tracking problems. Furthermore, aspects of TARSIA such as the symbolic/numeric interaction, are applicable to other domains where parameter estimation is performed in an uncertain or incompletely specified environment.

Acknowledgement

The authors would like to thank Mr. Kai Gong for his insight and encouragement in this project.

References

- [1] Gong, K. F., Magliaro, A. A., and Baylog, J. G., "A Decision-Directed Approach to Solution Integration for Tracking in an Underwater Environment," Proceedings of the 18th Asilomar Conference on Circuits, Systems and Computers, November 1984.
- [2] Silva, A., and Beck, N., "AI Applied to Submarine Combat Control Systems," Proceedings of the Workshop on AI and Distributed Problem Solving, National Academy of Sciences, May 1985.
- [3] Hassab, J. C., "Contact Localization and Motion Analysis in the Ocean Environment: A Perspective," IEEE Journal of Ocean Engineering, Vol. OE-8, No. 3, July 1983.
- [4] Nardone, S. C., and Aidala, V. J., "Observability Criteria for Bearings-Only Target Motion Analysis," IEEE Transactions on Aerospace and Electronic Systems, Vol. AES-17, No. 2, March 1981.
- [5] Lindgren, A. G., Graham, M. L., and Gong, K. F., "Nonlinear Parameter Estimation with Segmented Data: Trajectory Estimation with Biased Measurements," Proceedings of the 19th Asilomar Conference on Circuits, Systems and Computers, November 1985.
- [6] Baylog, J. G., Magliaro, A. A., Gong, K. F., and Hammel, S. E., "Performance of Constrained Estimation for Tracking with a Linear Array," Proceedings of the 18th Asilomar Conference on Circuits, Systems and Computers, November 1984.
- [7] Hammel, S. E., Aidala, V. J., Gong, K. F., Lindgren, A. G., "Recursive versus Batch Processing Algorithms for Bearings-Only Tracking," Proceedings of Oceans '83, Vol. 1, August 1983.
- [8] Garvey, T., Lowrance, J., and Fischler, M., "An Inferencing Technique for Integrating Knowledge from Disparate Sources," Proceedings of the 7th International Conference on Artificial Intelligence, 1981.
- [9] Kahn, K., Gorry, G. A., "Mechanizing Temporal Knowledge," Artificial Intelligence -- An International Journal, Volume 9, Number 1, 1977.
- [10] Allen, J., "Towards a General Theory of Action and Time," Artificial Intelligence -- An International Journal, Volume 23, Number 2, 1984.
- [11] Vilain, M., "A System for Reasoning About Time," AAAI-1982 Conference Proceedings, 1982.
- [12] Allen, J., "An Interval-Based Representation of Temporal Knowledge," Proceedings of the 7th International Conference on Artificial Intelligence, 1981.
- [13] Eddy, W., and Pei, G., "Structure of Rule-Based Belief Functions," Technical Report No. 347, Carnegie-Mellon University, 1985.
- [14] Morik, K. and Rollinger, C., "The Real Estate Agent - Modeling Users by Uncertain Reasoning," Artificial Intelligence Magazine, Vol. 6, No. 2, Summer 1985.